

ABSTRACT

A method and apparatus are disclosed for sharing execution capacity among tasks executing in a real-time computing system. The present invention extends RMA techniques for characterizing system timing behavior and designing real-time systems. A high priority task having hard deadlines is paired with a lower priority task having soft deadlines. During an overload condition, the higher priority task can dynamically borrow execution time from the execution capacity of the lower priority task without affecting the schedulability of the rest of the system. The higher priority task is bolstered in a proportion to the capacity borrowed from the lower priority task, so that the combined utilization of the two tasks remains constant. The period of the degraded task is increased to compensate for the execution time that was loaned to the higher priority task. In addition, the priority of the lower priority task is modified to match the new period.

1100-19.app